MACROECONOMIC EFFECTS OF THE GENDER REVOLUTION*

DRAGO BERGHOLT†, LUCA FOSSO‡ AND FRANCESCO FURLANETTO§

September 2023 (Preliminary Draft)

Abstract: Postwar U.S. labor market data feature a substantial secular decline in employment and wage gaps between males and females. We set out to identify the underlying, structural drivers of these trends, and quantify their macroeconomic effects. For this purpose, we propose a novel time series model which estimates empirical trends in macro data, and then decomposes these trends into selected (aggregate and gender-specific) structural trends. Identification is achieved with restrictions from a neoclassical model featuring gender-specific labor. Our empirical results point to a secular rise in demand for female labor–which in the model is driven by female-biased labor productivity growth–as the dominant source of closing gender gaps in the U.S. labor market. Moreover, this structural trend has been quantitatively important for the persistent components in aggregate employment and GDP, accounting for about one-third of overall economic growth in the postwar U.S. economy. Finally, to understand gender-specific labor supply we find it crucial to account for skills: a secular rise in the supply of skilled females has largely been counteracted by a contraction in the supply of unskilled females, thus, explaining why trends in gender-specific supply do not show up in aggregate data.

Keywords: *Gender Inequality, Gender Gaps, VAR with Common Trends.*
JEL Classification: *C32, E10, E13, J2.*

*This work should not be reported as representing the views of the European Central Bank or Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of the European Central Bank or Norges Bank. We thank for useful comments Stefania Albanesi, Guido Ascari, Martin Beraja, Paolo Bonomolo, Fabio Canova, Efrem Castelnuovo, Marco Del Negro, Domenico Giannone, Joseba Martinez, Karel Mertens, Silvia Miranda Agrippino, Evi Pappa, Giorgio Primiceri, Giuseppe Ragusa, Giovanni Ricco, Juan Rubio-Ramirez, Aysegul Sahin, Stefan Schiman, seminar participants at Norges Bank, Chicago Fed, Dallas Fed, University of Aix-Marseille, 2021 Sailing the Macro Conference in Ventotene, 2021 conference in honor of Fabio Canova in Hydra, 2022 Padova Macro Talks, 2022 Barcelona Summer Forum, 2022 Dolomiti Macro Meetings in Selva Val Gardena, University of Milano-Bicocca, 2022 Forskermøtet in Stavanger, 2023 SNDE conference in Orlando, 2023 IAAE conference in Oslo, 2023 CEF conference in Nice and forthcoming presentations at 2023 Dynare Conference in Malta, the University of Texas Austin, University of California Davis, San Francisco Fed, 2023 Fall Midwest Macro Meetings in Lubbock, Cleveland Fed and Washington University of Saint Louis.
†Norges Bank. Email: drago.bergholt@norges-bank.no
‡European Central Bank. Email: luca.fosso@ecb.europa.eu
§Norges Bank and BI Norwegian Business School. Email: francesco.furlanetto@norges-bank.no

1 Introduction

Women's increased labor market participation is arguably one of the fundamental changes observed in modern economies during the last century. Consider, for example, the U.S. labor market data in Figure 1: the employment rate for females in the U.S. in 1960 was less than half of the employment rate for males. But the ratio of female to male employment, or *the female employment gap*, increased steadily to 70 percent in the mid-1980s, before converging more gradually to around 85 percent in recent decades. The gender difference in wages displays a similar trend: women's hourly wages relative to men's wages, or *the female wage gap*, stayed relatively flat at 60 percent until the mid-1970s (despite a substantial catch-up in female employment rates during that period). But since then it has grown at about the same pace as the employment gap on average, resulting in a major wage convergence between females and males over time. In total, the wage gap has shrunk by about 50 percent, and the employment gap by more than 60 percent, over the last 5-6 decades. It is hardly a coincidence that Goldin (2006) talks about a "quiet revolution" when describing these labor market trends.

The goal of our paper is to investigate the macroeconomic consequences of the gender revolution. More specifically, we quantify its impact on economic growth in the US in terms of GDP, employment, and various measures of productivity. In fact, one can think that talent was highly misallocated in an economy in which only 45 percent of females were working, as shown in Hsieh, Hurst, Jones, and Klenow (2019). In addition, our second question of interest is to investigate what factors lie behind the gender convergence in employment and wages. The fact that these two trends co-move seems to suggest that labor demand factors must be dominant. However, it is undeniable that a large shock to the supply of skills, and thus a labor supply factor, has also affected the US economy. Our aim is to appropriately disentangle labor demand and labor supply factors once we take into account that a large increase in employment of female workers was concentrated in the market for high-skilled workers.

In order to answer our research questions, we use neoclassical macroeconomic theory in combination with a Structural Vector Autoregressive (SVAR) model, arguably the most standard tool for time series analysis, estimated on macroeconomic data and on gender-specific variables obtained by combining CPS data in the gender dimension with data from Dolado, Motyovszki, and Pappa (2021) to account for the skill dimension. Usually, SVAR models are mainly used to study cyclical fluctuations. In contrast, one key aspect of the question at hand is the focus on the macroeconomic effects of slow moving trends, thus implying that SVAR models need to be amended substantially. In practice, we use a SVAR model with common trends as in Del Negro, Giannone, Giannoni, and Tambalotti (2017) and Crump, Eusepi, Giannoni, and Sahin (2019). The model can be seen as a multivariate unobserved component model in which the variables enter in levels and in which our interest is on the permanent component, and not on the cyclical component as it is often the case in the literature. Put differently, rather than focusing on structural shocks driving the cyclical component, we conduct structural analysis on the permanent component and decompose it into various structural components, both aggregate and gender-specific. As an example, let us focus on GDP. Our model decomposes GDP dynamics into a cyclical component and a permanent component which, in turn, is driven by (gender neutral) technology shocks, (gender neutral) automation shocks, (gen-

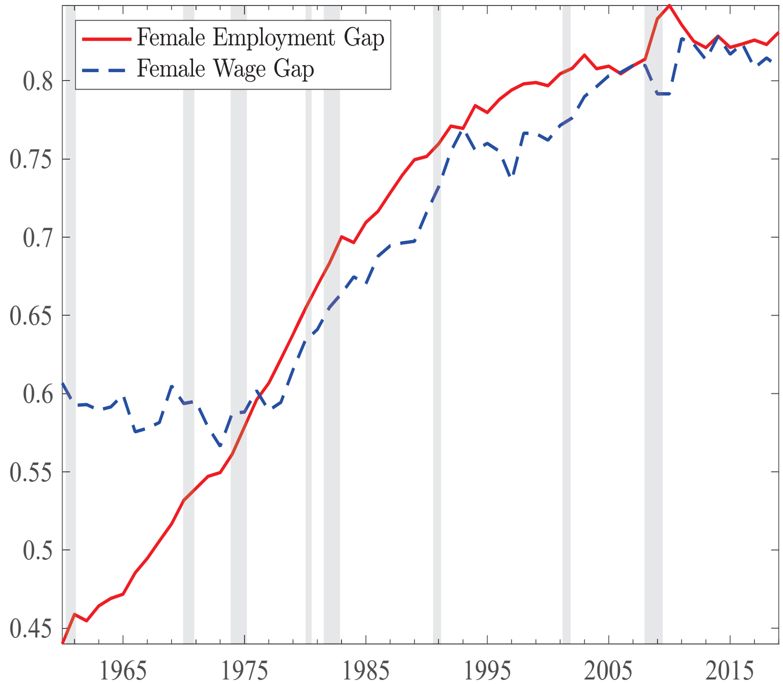

Figure 1: The female employment gap is the ratio between female employment and male employment, both measured relative to their respective populations. The female wage gap is the ratio between hourly female wages and male wages.

der neutral) labor supply shocks, gender-specific labor demand shocks and gender-specific labor supply shocks. Naturally, it is crucial to identify these structural drivers. Our key contribution is to derive identifying restrictions from a neoclassical model with gender which allows to disentangle the five *structural* driving forces based on their long-run impact on macroeconomic variables. This represents a key distinction from previous studies estimating VARs with stochastic trends. Del Negro et al. (2017) and Crump et al. (2019), for instance, allow for the presence of common trends in the data, but remain silent about the underlying *structural* sources of these slow-moving dynamics. Our framework, instead, allows to make a step further and identify the structural trends, by instructing the Bayesian algorithm with the prior information inherited from theory.

Our main result is that the forces driving the gender convergence in employment and wages are important also for trend GDP growth in the US. They account for up to a half between 1960 and 2000 while their contribution slows down to 15 per cent during the last 20 years, in keeping with the flattening of the employment gap dynamics. In addition, gender specific forces explain also a substantial share of aggregate employment dynamics over the entire sample period. When it comes to the individual role of the two gender specific forces, our model, perhaps not surprisingly, attributes a dominant role to gender-specific labor demand factors. However, it is too premature to declare that gender-specific labor supply factors are irrelevant. In a second step, in fact, we re-estimate our model by using data on the gender employment gap and on the gender wage gap within skilled workers. In the context of this more disaggregate exercise, we confirm that gender shocks are important for US economic growth. In this case, however, the gender employment gap is driven by both a gender labor demand shock and a gender supply shock. Both shocks are quantitatively important. Thus, using data disaggregated by skill allow us to identify an important shock to the supply of skills that was almost irrelevant in the baseline model. Why does such a shock emerge? The gender employment gap within skilled worker has converged faster than the aggregate gender employment gap. At the same time, the gender wage gap within skilled workers has seen a slower convergence than its aggregate counterpart, a fact documented also by Taniguchi and Yamada (2020). Both features are consistent with an important role for a shock to the labor supply of skilled female workers. In contrast, when we focus only on unskilled workers, we document that the gender employment ratio has barely moved over the last 60 years in the US while the gender wage gap has converged significantly faster within unskilled workers than within skilled workers. When we estimate our model using data on gender employment gap and gender wage gap only for unskilled workers, we find that a negative shock to the supply of unskilled workers is needed to reconcile the absence of convergence in employment and the strong convergence in wages between males and females. All in all, it seems that our baseline finds no role for the gender specific labor supply shock because a positive shock to the supply of skilled female workers is compensated by a negative shock to the supply of unskilled female workers. We conclude that labor supply factors are in fact important once the skill dimension in taken into account.

Our paper contributes to two strands of the literature. First, we contribute to a large literature studying the gender revolution. A useful distinction for our purposes is between papers emphasizing labor demand factors from labor supply factors. Among the former, Galor and Weil (1996) and Buera and Kaboski (2012) emphasize technological factors that favored the demand for women in combination with an increase in the returns to intellectual skills and the rise of the service sector, while Hsieh et al. (2019) point to a reduction in gender discrimination as an important driver of the reduction in the gender wage gap. Among the latter, Albanesi and Olivetti (2016) and Goldin and Katz (2002) document the importance of advances in maternal health and contraception, Fernández, Fogli, and Olivetti (2004) emphasize the importance of cultural factors developed during World War II, Attanasio, Low, and Sánchez-Marcos (2008) point to the crucial role of availability and affordability of child care, while Greenwood, Seshadri, and Yorukoglu (2005) propose a model in which the emergence of home appliances favors female's market production at the expense of home production. We contribute to this literature by proposing a horse race between labor supply and labor demand factors in the context of a macroeconomic time-series model. While less detailed in terms of the underlying transmission mechanisms, our analysis provide a clear link between gender trends and macroeconomic outcomes.

We contribute also to a recent literature emphasizing the role of gender for macroeconomic dynamics in quantitative set-ups. Heathcote, Storesletten, and Violante (2017) and Hsieh et al. (2019) propose a decomposition of US macroeconomic growth in structural models with gender. Albanesi (2019) estimates with Bayesian methods a real business cycle model with gender with a focus on the importance of gender trends to account for jobless recoveries, while Fukui, Nakamura, and Steinsson (2023) propose a similar model with a focus on the fact that rising female participation has not crowded out male participation and has thus been an expansionary factor for the US economy. In a similar spirit to Heathcote et al. (2017) and Hsieh et al. (2019), we propose a decomposition of US trend economic growth but, differently from their approach, within a SVAR framework, whose identification is motivated by theory.

Finally, the remainder of the paper is organised as follows. Section 2 presents the empirical model. The analytical solutions from the neoclassical model are derived in section 3. Section 4 discusses the identification strategy. Sections 5 and 6 are devoted to results. Concluding remarks are exposed in section 7.

2 A TIME SERIES MODEL WITH COMMON TRENDS

The model that we estimate is a multivariate time series model with unobserved components. It is designed to pin down two objects simultaneously: first, it decomposes a vector of observable data into unobservable cycle and trend components. Second, it further maps the empirical trends into a vector of structural trends which are stochastic. The relationship between empirical and structural trends is assumed linear, but can otherwise be arbitrarily flexible. This implies that trends in observable data may share common components–the underlying structural drivers. To fix ideas, consider an $n \times 1$ vector of data Y_t, which is the sum of two unobserved states:

$$Y_t = \hat{Y}_t + \bar{Y}_t, \tag{1}$$

$\hat{Y}_t$ and $\bar{Y}_t$ represent the empirical cycle and trend, respectively. Equation (1) is a purely *statistical* decomposition of data. The main focus of our analysis will be on the trend block of this decomposition and, more precisely, on the underlying *structural* forces behind $\bar{Y}_t$. Suppose that there are $q \leq n$ structural forces at play:

$$\bar{Y}_t = \mathcal{V} X_t \tag{2}$$

X_t is the $q \times 1$ vector of structural trends and $\mathcal{V}$ is the $n \times q$ matrix that maps the reduced-form trends into structural ones. Importantly, $\mathcal{V}$ embeds the long-run identifying restrictions required to uniquely pin down the structural trends and reconcile X_t with $\bar{Y}_t$. Similarly to Del Negro et al. (2017) and Crump et al. (2019), we assume that each of the structural trends follows a random walk, potentially with a drift:[1]

$$X_t = c + X_{t-1} + u_t, \qquad u_t \sim \mathcal{N}(0_q, \Sigma_u) \tag{3}$$

Throughout we assume that the covariance matrix Σ_u is diagonal. Since our focus is on trends rather than the cyclical part of data, only a minimal set of restrictions is imposed on $\hat{Y}_t$. In particular, we model $\hat{Y}_t$ as a stationary, reduced-form vector autoregressive (VAR) process:

$$\Phi(L)\hat{Y}_t = e_t, \qquad e_t \sim \mathcal{N}(0_n, \Sigma_e) \tag{4}$$

$\Phi(L) = I - \Phi_1 L - ... - \Phi_p L^p$ is an $n \times n$ matrix of lag coefficients. Σ_e is freely estimated without any restrictions on the off-diagonal elements. However, we assume that permanent and transitory shocks are mutually uncorrelated, i.e. that $cov(u_t, e_t) = 0$.[2]

[1] More recent extensions of the same model to explain inflation dynamics include Ascari and Fosso (2021), Bianchi, Nicolò, and Song (2023) and Maffei-Faccioli (2021).

[2] This suggests, for instance, that by construction a technology trend shock does not affect the cycle, as instead conventional in the RBC literature. Though being a strong assumption, it should not be of great

Equations (1)-(4) constitute the model that we confront with data. For reasons that will be clear below, we include data on aggregate GDP, wages and employment, as well as data on female-to-male differences in wages and employment, respectively. The latter two variables allow us to identify a structural trend in female-specific labor productivity, as well as a structural trend in female-specific labor supply. One essential object of interest is the matrix $\mathcal{V}$ which determines how, and to what extent, the structural factors in X_t give rise to common trends $\bar{Y}_t$ in data. A key methodological contribution of this paper is that we estimate parts of $\mathcal{V}$, using economic theory to achieve identification. The theoretical framework is presented next.

3 A STYLIZED MODEL WITH GENDER-SPECIFIC LABOR

In this section we present a neoclassical model that serves the purpose of deriving theory-based identification assumptions and prior distributions in order to estimate the matrix $\mathcal{V}$. The theoretical model builds on and extends the previous key contributions by Fukui et al. (2023) and Albanesi (2019).

The model economy is populated by a unit mass of identical firms, and a unit mass of identical households who own equal shares in the firms. A representative firm chooses labor inputs and capital investments in order to maximize a properly discounted sum of expected lifetime profits, $\mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \frac{\Lambda_s}{\Lambda_t} \Pi_s$. For each period t we denote the rational expectations operator (conditional on the information currently available) by $\mathbb{E}_t$. $\beta \mathbb{E}_t \frac{\Lambda_s}{\Lambda_t}$ captures households' discounting of the future where β is the time discount rate and Λ_t represents the shadow value of income. The firm's period profit is equal to

$$\Pi_t = Y_t - W_{f,t} L_{f,t} - W_{m,t} L_{m,t} - P_{I,t} I_t. \tag{5}$$

Y_t represents output, $W_{f,t}$ ($W_{m,t}$) represents the real wage rate specific to female (male) labor, and $L_{f,t}$ ($L_{m,t}$) is the quantity of female (male) labor used in production. I_t represents gross investments in physical capital. The relative price of investments is given by $P_{I,t}$. The firm's maximization problem is subject to the production function

$$Y_t = A_t L_t^{\alpha_t} K_{t-1}^{1-\alpha_t}, \tag{6}$$

where K_{t-1} stands for physical capital currently in place, and L_t is an aggregation of male and female labor:

$$L_t = \left[\alpha_l \left(A_{m,t} L_{m,t} \right)^{\frac{\gamma-1}{\gamma}} + (1-\alpha_l) \left(A_{f,t} L_{f,t} \right)^{\frac{\gamma-1}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}} \tag{7}$$

A_t, $A_{m,t}$ and $A_{f,t}$ are aggregate and gender-biased productivity shocks, respectively, while $\gamma > 1$ governs the degree of substitution between genders when firms produce. Note that we allow for a time varying weight α_t on aggregate labor. One possible interpretation

concern in our case, provided that we are solely interested in modeling the secular trends dynamics. Relatedly, one may wonder how our model interpret business cycle shocks that that leave long-run traces (hysteresis effects) on the productivity capacity of the economy. Any shock with long-run effects is captured by the permanent component in our model. We do not identify a separate trend driven by demand factors because the evidence in favor of hysteresis effects is confined to the last 40 years (Furlanetto, Lepetit, Ørjan Robstad, Rubio-Ramírez, and Ulvedal (2023)) while our sample starts well before in 1960.

of a decline in α_t is that it follows from labor-displacing automation, see Acemoglu and Restrepo (2020) and Bergholt, Furlanetto, and Maffei-Faccioli (2022). Finally, physical capital dynamics are given by

$$K_t = (1-\delta) K_{t-1} + I_t. \tag{8}$$

The representative firm's first order conditions with respect to investments, capital, and male and female labor, are summarized below:

$$P_{I,t} = \mathcal{Q}_t \tag{9}$$

$$\mathcal{Q}_t = \beta \mathbb{E}_t \frac{\Lambda_{t+1}}{\Lambda_t} \left[(1-\alpha_{t+1}) \frac{Y_{t+1}}{K_t} + \mathcal{Q}_{t+1} (1-\delta) \right] \tag{10}$$

$$W_{m,t} = \alpha_t \alpha_l \frac{Y_t}{L_t} \left(\frac{L_t}{L_{m,t}} \right)^{\frac{1}{\gamma}} A_{m,t}^{\frac{\gamma-1}{\gamma}} \tag{11}$$

$$W_{f,t} = \alpha_t (1-\alpha_l) \frac{Y_t}{L_t} \left(\frac{L_t}{L_{f,t}} \right)^{\frac{1}{\gamma}} A_{f,t}^{\frac{\gamma-1}{\gamma}} \tag{12}$$

The first optimality condition states that firms invest until the price of investment is equal to $\mathcal{Q}_t$, the shadow value of one more unit of installed capital in the next period. The second optimality condition defines the shadow value of capital: it is the properly discounted sum of next period's marginal product of capital and the continuation value net of depreciation. The two last optimality conditions pin down optimal firm demand for male and female labor, respectively. Everything else equal, gender-specific labor demand is increasing in aggregate activity, decreasing in the gender-specific wage rate, and, if $\gamma > 1$, increasing in gender-specific productivity.

The representative household is populated by an equal number of male and female workers. In each period it chooses a plan for consumption and labor supply in order to maximize expected lifetime welfare $\mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \mathcal{U}_s$, where

$$\mathcal{U}_t = \frac{C_t^{1-\sigma}}{1-\sigma} \exp\left(-\Psi_t^{-1} \frac{(1-\sigma) \tilde{L}_t^{1+\varphi}}{1+\varphi} \right) \tag{13}$$

represents the period utility function. Aggregate labor dis-utility $\tilde{L}_t$ is increasing in male and female labor:

$$\tilde{L}_t = \left[\left(\frac{L_{m,t}}{\Psi_{m,t}} \right)^{\frac{1+\lambda}{\lambda}} + \left(\frac{L_{f,t}}{\Psi_{f,t}} \right)^{\frac{1+\lambda}{\lambda}} \right]^{\frac{\lambda}{1+\lambda}} \tag{14}$$

$\Psi_{m,t}$ and $\Psi_{f,t}$ are gender-specific labor supply shocks, $\lambda > 0$ governs the household's willingness to substitute female with male labor. The representative household's first order conditions with respect to consumption, bond savings, as well as supply of male and female labor respectively, are summarized below:

$$\Lambda_t = C_t^{-\sigma} \exp\left(-\Psi_t^{-1} \frac{(1-\sigma) \tilde{L}_t^{1+\varphi}}{1+\varphi} \right) \tag{15}$$

$$\Lambda_t = \beta \mathbb{E}_t \Lambda_{t+1} (1+r_t) \tag{16}$$

$$W_{m,t} = \Psi_t^{-1} C_t \tilde{L}_t^{\varphi - \frac{1}{\lambda}} L_{m,t}^{\frac{1}{\lambda}} \Psi_{m,t}^{-\frac{1+\lambda}{\lambda}} \tag{17}$$

$$W_{f,t} = \Psi_t^{-1} C_t \tilde{L}_t^{\varphi - \frac{1}{\lambda}} L_{f,t}^{\frac{1}{\lambda}} \Psi_{f,t}^{-\frac{1+\lambda}{\lambda}} \tag{18}$$

The first optimality condition equates the shadow value of income with the marginal utility of consumption. The second optimality condition states the optimal, intertemporal consumption plan. The two last optimality conditions illustrate that, everything else equal, the optimal supply of gender-specific labor is increasing in the gender-specific wage rate, decreasing in aggregate consumption, and increasing in the aggregate and gender-specific labor supply shocks. Finally, an increase in male labor for example, which in turn raises aggregate labor dis-utility $\tilde{L}_t$, implies a reduction (increase) in female labor supply if and only if $\varphi > \lambda^{-1}$ ($< \lambda^{-1}$). Importantly, λ governs the household's willingness to substitute work across genders. A sufficiently low value of λ implies complementarity of gender-specific labor dis-utility to such an extent that more time spent working for the male causes a decline in the joy of leisure for the female.

In order to characterize gender differences in the labor market, we find it instructive to focus on the female wage gap $w_{f,t} = \frac{W_{f,t}}{W_{m,t}}$, as well as the female employment gap $l_{f,t} = \frac{L_{f,t}}{L_{m,t}}$. This notation allows us to combine the firm's optimality conditions with respect to male and female labor in order to express relative labor demand, which is downward sloping in the $(w_{f,t}, l_{f,t})$-space:

$$l_{f,t} = \left(\frac{1-\alpha_l}{\alpha_l} \right)^{\gamma} w_{f,t}^{-\gamma} a_{f,t}^{\gamma-1} \tag{19}$$

The slope coefficient $-\gamma$ determines how responsive demand is to relative wage changes. It follows naturally that shifts in $l_{f,t}$ not associated with changes in $w_{f,t}$ are driven by the "ratio shock" $a_{f,t} = \frac{A_{f,t}}{A_{m,t}}$, which we interpret as a relative demand shifter. In a similar way, we can combine the household's optimality conditions with respect to male and female labor in order to express relative labor supply, which is sloping upwards in the $(w_{f,t}, l_{f,t})$-space. The slope coefficient λ determines how responsive supply is to relative wage changes:

$$l_{f,t} = w_{f,t}^{\lambda} \psi_{f,t}^{1+\lambda} \tag{20}$$

We interpret the "ratio shock" $\psi_{f,t} = \frac{\Psi_{f,t}}{\Psi_{m,t}}$ as a supply shifter which effectively soaks up all the variation in relative labor supply not associated with movements in the wage gap. Combining the two previous equations, one arrives at the following analytical solutions for the wage and employment gaps between females and males:

$$w_{f,t} = \left(\frac{1-\alpha_l}{\alpha_l} \right)^{\frac{\gamma}{\gamma+\lambda}} a_{f,t}^{\frac{\gamma-1}{\gamma+\lambda}} \psi_{f,t}^{-\frac{1+\lambda}{\gamma+\lambda}} \tag{21}$$

$$l_{f,t} = \left(\frac{1-\alpha_l}{\alpha_l} \right)^{\frac{\gamma\lambda}{\gamma+\lambda}} a_{f,t}^{\frac{(\gamma-1)\lambda}{\gamma+\lambda}} \psi_{f,t}^{\frac{(1+\lambda)\gamma}{\gamma+\lambda}} \tag{22}$$

Importantly, the female biased demand shock $a_{f,t}$ implies co-movement between wage and employment gaps across genders, while the female biased supply shock $\psi_{f,t}$ implies negative co-movement. Moreover, while macroeconomic shocks may drive the *absolute* *level* of female wages and employment, only the two "ratio shocks" $a_{f,t}$ and $\psi_{f,t}$ can affect $w_{f,t}$ and $l_{f,t}$, i.e. the *relative* wage and employment levels of female workers. "Macro shocks" such as A_t and Ψ_t play no role here.[3] A corollary statement is that gender-specific labor market variables and their aggregate counterparts display proportional responses to macroeconomic shocks (e.g. $L_{f,t} \propto L_t$). Importantly, these model implications form a key part of our identification scheme in the empirical section, allowing us to disentangle the different structural drivers of $w_{f,t}$ and $l_{f,t}$ in data.

Finally, we also note that the model presented here nests as special cases some recent theoretical contributions in the literature. Fukui et al. (2023), for example, implicitly assume $\gamma = \infty$, i.e. perfect gender substitutability within the firm. By construction this causes the wage gap to be driven solely by gender biased demand shocks, as can be seen from the analytical solution for $w_{f,t}$ given above. Albanesi (2019), in contrast, implicitly assumes that $\lambda = \frac{1}{\phi}$. This knife-edge parametrization effectively makes the supply of female and male labor independent of how much the spouse is working, as emphasized earlier. While our estimation procedure allows for these special cases, a potentially important contribution of our paper is to quantify the degree of gender complementarity–both on the firm and household side–in the labor market. This completes our description of the theoretical framework.

4 From theory to trend identification in data

Next we explain how the neoclassical theory just described is used in practice to discipline our empirical analysis, allowing us to identify the underlying, structural trends in data.

4.1 Monte Carlo simulations

The main purpose of the theoretical model is to infer a set of theory-consistent restrictions on $\mathcal{V}$. These restrictions should be (i) informative enough so that they ensure the identification of all estimated elements in $\mathcal{V}$, and (ii) sufficiently agnostic so that our analysis remains robust to all reasonable parametric perturbations of the underlying theory.

To this end we conduct the following simulation exercise: first, we draw a parameter vector $\theta = [\sigma, \varphi, \gamma, \lambda, ...]'$ which includes all parameters of interest in the theoretical model, including state variables such as initial wage and employment gaps. In order to be as agnostic as possible, we draw each parameter independently from a uniform distribution specified further below. Second, conditional on θ, we solve the theoretical model numerically and compute the impulse responses of macro and gender variables to each of the structural trend shocks. For our purpose, the main interest lies in the long end of the impulse responses, i.e. the long-run effects of various trend shocks. For that reason, we compute the perfect foresight solution of the model. We repeat the exercise $1,000$ times and save all impulse responses. This Monte Carlo exercise leaves us, at each time horizon and conditional on each shock, with an entire distribution of structural outcomes for the endogenous variables of interest. The distribution visualizes variation in outcomes due to parameter uncertainty.

[3] This is true not only in the long run but also within the business cycle. The irrelevance of aggregate macro shocks for gender gaps is a consequence of the constancy of gender substitution elasticities γ and λ, and remains even if we were to introduce business cycle frictions such as nominal price rigidities.

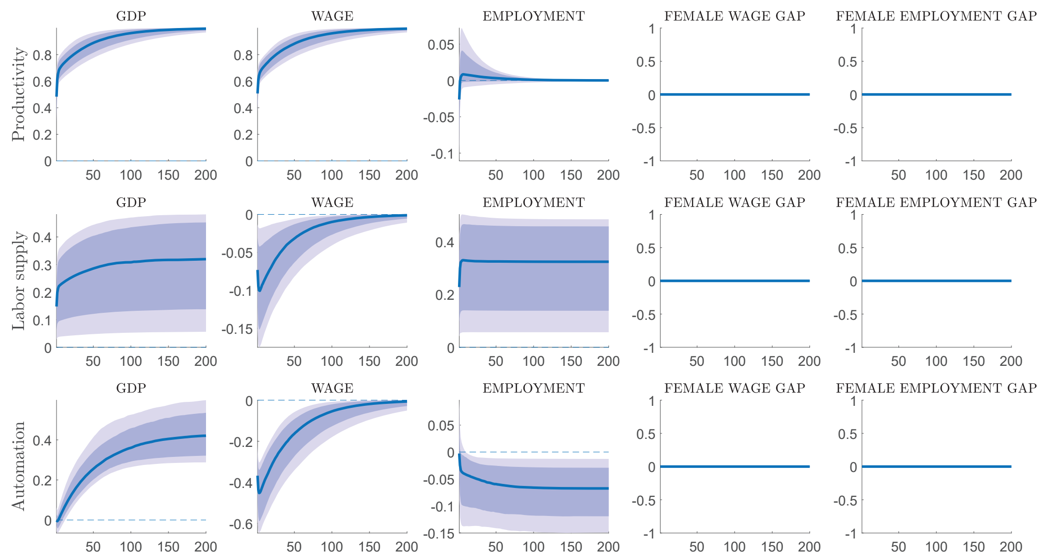

Figure 2: Impulse responses to aggregate macro shocks.

Regarding the uniform distributions for the structural parameters, we impose bounds that are wide enough so that they span the set of values proposed in existing literature. In particular, for the three "macro" parameters σ, φ and α we choose $\sigma \sim U(1, 5)$, $\varphi \sim U(0, 4)$ and $\alpha \sim U(0.5, 0.7)$ respectively. Note that common values for the aggregate Frisch elasticity φ^{-1}, both from microeconomic and macroeconomic research, are well within the bounds used here. Moreover, the bounds on α allow the model to cover a wide range of labor income shares, including all those observed in the postwar US economy. There is substantially less external information available about the key gender parameters of interest, γ and λ. For these we choose the following uniform distributions: $\gamma \sim U(1, 11)$ and $\lambda \sim U(0, 1)$. Regarding γ, we note that the median value is $\gamma \approx 5$, which is similar to the value proposed by Albanesi (2019). Fukui et al. (2023) in contrast assume $\gamma = \infty$ in their baseline, but also analyze the case with $\gamma = 5$ as a robustness test. Finally, note that $\varphi = \frac{1}{\lambda}$ for median values of φ and λ. This is exactly the special case considered by both Albanesi (2019) and Fukui et al. (2023). However, our chosen bounds allow us to investigate very different scenarios as well, including those with substantial complementarity (as well as substitutability) across genders when the household decides on labor supply. When solving the model, the initial wage and employment gaps may matter for the long run outcomes of structural shocks. Therefore, we choose to draw initial wage and employment gaps from $w_{f,0} \sim U(0.56, 0.85)$ and $l_{f,0} \sim U(0.44, 0.85)$ respectively. These bounds are chosen so that they cover both the highest and the lowest wage and employment gaps observed in the postwar US economy.

Impulse responses from the simulation exercise are presented in figures 2 and 3. Since the goal is to disentangle structural long-run trends, we restrict our attention to the long end of the impulse responses. We start by assessing the responses to aggregate shocks in figure 2. A few remarks are in place: first, consistent with the analytical solutions in equations (21)-(22), aggregate shocks have no effects whatsoever on the wage and employment gaps. This is true at all horizons. Second, conditional on an aggregate productivity shock, GDP and aggregate wages display the same response both in the sign and the magnitude,

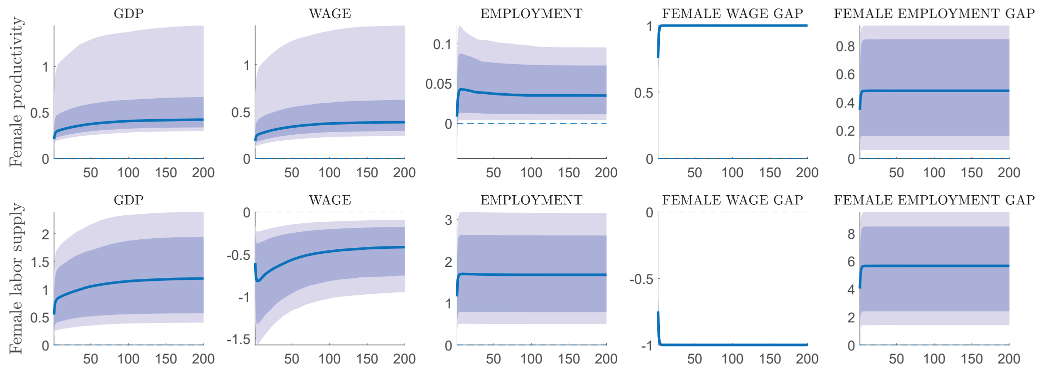

Figure 3: Impulse responses to gender-specific shocks.

while aggregate employment is not affected in the long run. Third, in response to an aggregate labor supply shock, we obtain long-run co-movement between GDP and aggregate employment, but no long-run effects on aggregate wages. Finally, an automation shock generates negative co-movement between GDP and aggregate employment, but has no long-run effects on aggregate wages.

Let us now move to the gender-specific trend shocks in figure 3. Both shocks are normalized to produce a unit effect on the wage gap in the long run.[4] Conditional on a gender-specific productivity shock, or more precisely a permanent rise in female-specific productivity, both the wage gap and the employment gap display a *positive* response. Since this shock increases average labor productivity in the economy, all three aggregate macro variables rise as well in the long run. In contrast, a gender-specific labor supply shock, or more precisely a permanent fall in females' labor dis-utility, causes the wage and employment gaps to respond in opposite directions. Moreover, GDP and employment both increase in the long run while the aggregate wage rate falls.

Overall, these simulation results provide us with a set of theory-consistent identification restrictions that we can impose on the matrix $\mathcal{V}$ when the empirical model is estimated on data. The restrictions enable us to simultaneously identify (i) three aggregate trends that characterize the long-run behavior of aggregate GDP, employment and wages, but at the same time have zero effect on gender differences in wages and employment, (ii) a gender-specific labor demand trend that causes the wage gap and the employment gap to go in the same direction, and (iii) a gender-specific labor supply trend causing the wage and employment gaps to go in opposite directions.

4.2 REVISITING THE MAPPING TO EMPIRICAL TRENDS

Recall equation (2), which links empirical trends to the underlying structural drivers. Given the results from the Monte Carlo simulations discussed above, we are now in a position to specify a baseline version of the vector of empirical trends $\bar{Y}_t$, the vector of

[4]This normalization is particularly useful when we later estimate the empirical model, as it will enable us to interpret the long-run effects of these two shocks on the employment gap in terms of the elasticities γ and λ.

structural trends X_t, and the mapping from structural to empirical trends $\mathcal{V}$:

$$\underbrace{\begin{bmatrix} G\bar{D}P_t \\ \bar{W}_t \\ \bar{E}_t \\ \bar{W}_{f-m,t} \\ \bar{E}_{f-m,t} \end{bmatrix}}_{\bar{Y}_t} = \underbrace{\begin{bmatrix} 1 & 1 & 1 & \nu_{14} & \nu_{15} \\ 1 & 0 & 0 & \nu_{24} & \nu_{25} \\ 0 & \nu_{32} & 1 & \nu_{34} & \nu_{35} \\ 0 & 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & \gamma & \lambda \end{bmatrix}}_{\mathcal{V}} \underbrace{\begin{bmatrix} A_t \\ \mathcal{M}_t \\ \Psi_t \\ \psi_{f,t} \\ a_{f,t} \end{bmatrix}}_{X_t} \tag{23}$$

The first column of $\mathcal{V}$ imposes restrictions on the stochastic technology trend A_t. Consistent with the Monte Carlo results summarized in Figure 2, we assume that a unit rise in the long-run technology level implies a unit rise in long-run GDP and aggregate wages, $G\bar{D}P_t$ and $\bar{W}_t$, but that it has a zero long-run effect on aggregate employment $\bar{E}_t$, as well as zero long-run effects on female-to-male gaps in wages and employment, $\bar{W}_{f,t}$ and $\bar{E}_{f,t}$. The second column of $\mathcal{V}$ governs the long-run effects of the stochastic trend in automation $\mathcal{M}_t$. Consistent with Figure 2 we impose that a permanent rise in automation has a positive effect on long-run GDP, normalized to one, a zero effect on long-run wages, as well as a negative effect on long-run employment. The latter restriction implies that $\nu_{32} < 0$. The gender gap trends are not affected by automation. The third column in $\mathcal{V}$ imposes restrictions on the aggregate labor supply trend Ψ_t. A rise in Ψ_t represents a permanent expansion of aggregate labor supply. Consistent with Figure 2, this implies upward shifts in trend GDP and trend employment. The trend in aggregate wages, as well as the gender gap trends, are not affected by the permanent labor supply shock.

The fourth and fifth columns in $\mathcal{V}$ govern long-run implications of the stochastic trends in female-specific labor supply and labor demand, respectively. Consistent with the Monte Carlo results summarized in Figure 3, we impose that a permanent rise in female-specific productivity causes the female-to-male gaps in wages and employment to co-move, while the opposite is the case for a permanent, expansionary shock to female-specific labor supply. Note that we normalize both of the gender trends so that they have a unitary effect on the gender gap in real wages. This choice turns out to be particularly convenient: rather than estimating $a_{f,t}$ and $\psi_{f,t}$, the normalization implies that we identify $\tilde{a}_{f,t} \equiv a_{f,t}^{\frac{\gamma-1}{\gamma+\lambda}}$ and $\tilde{\psi}_{f,t} \equiv \psi_{f,t}^{-\frac{1+\lambda}{\gamma+\lambda}}$. This allows us to specify the log-linearized versions of (21)-(22) as follows:

$$\begin{aligned} \hat{w}_{f,t} &= c_{w,f} + \hat{\tilde{a}}_{f,t} - \hat{\tilde{\psi}}_{f,t} \\ \hat{l}_{f,t} &= c_{l,f} + \lambda\hat{\tilde{a}}_{f,t} + \gamma\hat{\tilde{\psi}}_{f,t} \end{aligned}$$

A hat means that the variable is expressed in logarithms, with $c_{w,f}$ and $c_{l,f}$ being reduced-form constants. Importantly, the two gender elasticities λ and γ, two structural parameters of particular interest, enter directly as coefficients in the employment gap equation above. This means that $\lambda = \nu_{45}$ and $\gamma = \nu_{55}$ can be read directly from the estimated matrix $\mathcal{V}$. Finally, regarding the elasticities that govern feedback from gender trends to macroeconomic trends in GDP, wages and employment, we once again exploit the Monte Carlo results in Figure 3. The female-specific labor productivity shock for example behaves similarly to a conventional technology shock in the aggregate economy, with GDP and real wages rising permanently. Female-specific labor supply, in contrast, causes GDP and

		Prior		Posterior		
		Density	Support	Mean	Mode	90% HPD
ν_{14}	$a_f \rightarrow G\bar{D}P$	Uniform	$[0, 1]$	0.93	0.98	(0.85, 0.99)
ν_{24}	$a_f \rightarrow \bar{W}$	Uniform	$[0, 1]$	0.64	0.63	(0.45, 0.82)
ν_{34}	$a_f \rightarrow \bar{E}$	Uniform	$[-0.5, 0.5]$	0.32	0.34	(0.19, 0.45)
ν_{15}	$\psi_f \rightarrow G\bar{D}P$	Uniform	$[0, 2]$	0.24	0.06	(0.03, 0.53)
ν_{25}	$\psi_f \rightarrow \bar{W}$	Uniform	$[-2, 0]$	-0.17	-0.04	(-0.36, -0.02)
ν_{35}	$\psi_f \rightarrow \bar{E}$	Uniform	$[0, 3]$	0.52	0.34	(0.13, 1.11)
ν_{32}	$\mathcal{M} \rightarrow \bar{E}$	$\Gamma(6, .05)$	$(0, \infty)$	0.36	0.32	(0.23, 0.50)
λ	$a_f \rightarrow \bar{E}_{f-m,t}$	$\Gamma(2, .5)$	$(0, \infty)$	1.73	1.73	(1.39, 2.1)
γ	$\psi_f \rightarrow \bar{E}_{f-m,t}$	$\Gamma(6, .5)$	$(0, \infty)$	2.83	2.86	(1.78, 3.87)

Table 1: Prior distributions and posterior estimates. The posterior moments are generated from the last 10,000 of 50,000 draws generated from the RW Metropolis-Hastings algorithm.

employment to rise permanently while the aggregate real wage falls, just like an aggregate labor supply shock.

The restrictions imposed on each column in $\mathcal{V}$ are mutually exclusive, which is what we need to separately identify the five stochastic trends of interest. The aggregate technology trend, for example, is the only one that makes the real wage co-move with GDP in the long run. Automation and aggregate labor supply are separable because they imply opposite correlations between long-run wages and employment. Finally, female-specific labor supply and labor demand are separable from aggregate macro trends because they are the only drivers of long-run wage and employment gaps between females and males. Moreover, they are uniquely identified because they imply opposite signs on the co-movement between gender gaps in wages and employment.

4.3 Priors

Given the theory-consistent restrictions discussed above, we need to specify prior shapes for the estimated parameters. The priors used are summarized in Table 1. We aim for an agnostic approach and use uniform priors for all elasticities governing the feedback from gender trends to the aggregate macroeconomy. The support of these uniform priors largely reflects the uncertainty bands computed during the Monte Carlo exercise, as shown in Figure 2 and Figure 3 respectively. That is, consistent with theory, the female-specific productivity shock behaves qualitatively as a technology shock in the aggregate, while the female-specific labor supply shock behaves qualitatively as a gender-neutral labor supply shock. The prior for permanent automation effects on employment has a Gamma-distribution and is assumed negative, reflecting that automation crowds out employment.

The final two parameters that we estimate are γ and λ. The gender-specific labor demand elasticity γ has been quantified in a few existing studies. Weinberg (2000), for example, finds that γ is around 2.4 in the US, while Acemoglu, Autor, and Lyle (2004) report a slightly higher value of 3.[5] Thus, we choose a Gamma-prior for γ centered

[5]Johnson and Keane (2013) obtain estimates of γ spanning between 1.85 and 2.2 in a dynamic general

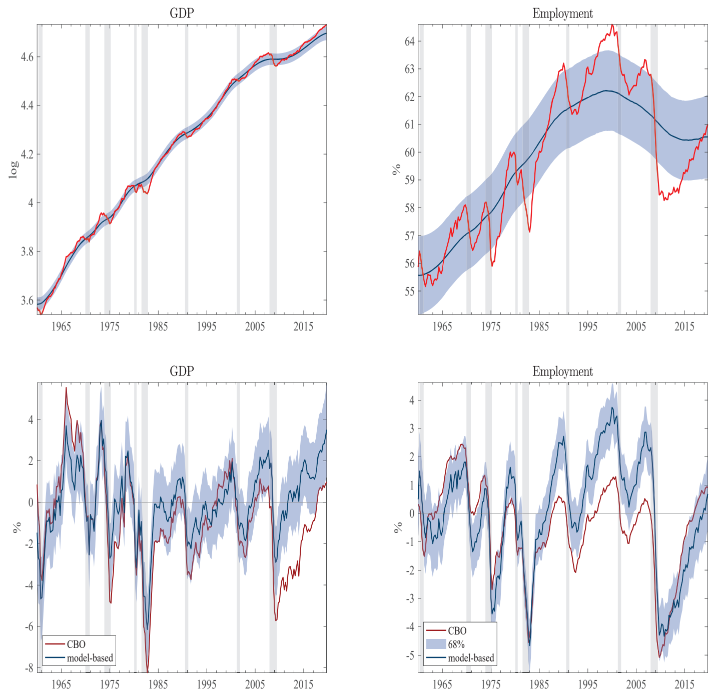

Figure 4: Estimated empirical trend and cycle of real GDP and aggregate employment. Left-hand side panel: observed data (red solid), median estimate (blue solid). Right-hand side panel: CBO gaps (dark red), median estimate (blue solid). 68% blue shaded areas.

around 3 with most of the probability mass located between 1 and 5. Literature on the gender-specific labor supply elasticity λ is scant, as far as we can tell. However, Keane and Rogerson (2012) show that relatively small (micro) labor supply elasticities can be reconciled with aggregate elasticities ranging between 1 and 2. In a review of the micro literature, in fact, Blundell and Macurdy (1999) conclude that λ must fall in a range of values below one. Thus, we choose a Gamma prior for λ with most mass below one, even though much higher values are allowed as well during estimation. Our setup captures quite well the specific case of Albanesi (2019) in which $\lambda = \frac{1}{\phi}$, given reasonable values for ϕ. Also, note that firms can switch between female and male labor more easily than households ($\lambda < \gamma$) at the prior mode, a feature that seems highly reasonable.

5 Results: baseline model

In this section we present results based on the time series model described in section 2. Given the theoretical restrictions derived in section 3, the vector of endogenous variables Y_t includes: (i) real GDP, (ii) real aggregate wages, (iii) the aggregate employment-to-population ratio, (iv) the ratio of female-to-male employment, and finally (v) the ratio of female-to-male wages. All variables enter the system in log-levels. Since we use

equilibrium model fitted to US data. Hamermesh (1993) reports values of approximately 2.3 and 2 for Australia and UK, respectively.

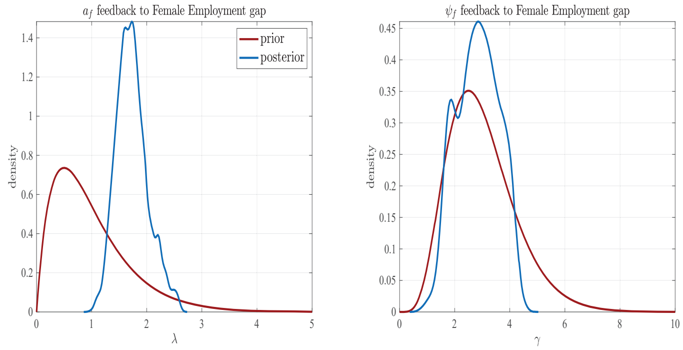

Figure 5: Prior distributions (red) and posterior distributions (blue) of the coefficients λ and γ.

quarterly data, the number of lags is chosen to be $p = 4$ in the baseline setup. The model is estimated over the sample 1960:Q1-2019:Q4.[6]

Permanent and transitory components. The first set of results is related to the decomposition into permanent and transitory components for all the observable variables. In the literature, the SVAR model with common trends has been used so far precisely to decompose data into a permanent component related to the "potential" or "natural" value of variables (like trend inflation in Ascari and Fosso (2021) or the natural rate of interest in Del Negro et al. (2017)) and a cyclical component (like the unemployment gap in Crump et al. (2019) or the inflation gap in Bianchi et al. (2023)). In our case, this is just a preliminary step to filter out the cyclical component given our focus on the structural drivers of the permanent component. In Figure 4 we present the estimated permanent component of output and employment together with the output gap and the employment gap, both defined as the difference between the observed series and its permanent component. It is interesting to see that the output gap tracks very well the estimate provided by the Congressional Budget Office (CBO) which constitutes a classic benchmark in the literature. The only noticeable difference is in the post-Great Recession period when there is a level shift between the two series. While this result is per se interesting (and related to Coibion, Gorodnichenko, and Ulate (2018)), it is almost irrelevant for our purposes given our interest in the gender convergence which happened way earlier in the sample. The same level shift appears in the employment gap before the Great Recession while the model tracks the CBO estimate very well in the post-Great Recession period. Note that such a similarity is neither obvious nor targeted since our SVAR is not informed by data on the CBO estimates. We conclude that the model offers a reasonable description of trend and cycles over US history and can be used as a laboratory to investigate the structural drivers of the permanent components. The same decomposition is proposed in the Appendix for the remaining variables (see Figure C.1).

Coefficients in $\mathcal{V}$. The estimated coefficients in the matrix $\mathcal{V}$ are a crucial input in order to perform our decomposition of the permanent components into various structural drivers. The posterior estimates of the nine estimated parameters are presented in Table 1.

The two most interesting estimates in $\mathcal{V}$ are presented separately in figure 5 where

[6]See Appendix 7 for the description and source of the data.

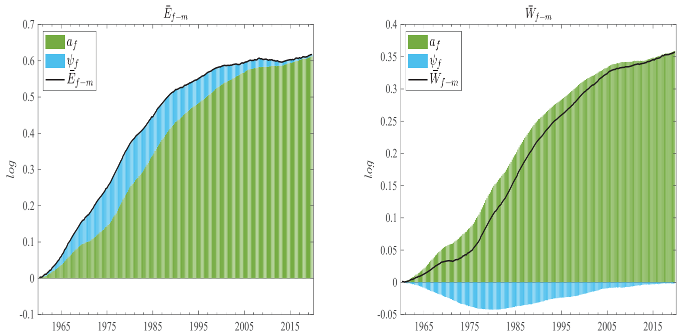

Figure 6: Contributions of the female-specific *structural* trends to the trends in female employment and wage gaps.

we show the posterior distributions of the elasticities λ and γ, the feedbacks of gender-specific labor demand and supply shocks to the female employment gap. Recall that these two parameters have a direct link with theory, as they describe the degree of female-male substitutability in the household and firm sectors of our theoretical model. In addition, it is important to stress once more that, differently from the other feedbacks to macro, the prior densities of λ and γ are based on the literature. This is indeed the reason why the priors (in red) are not distributed uniformly in the probability space. Despite a rather conservative on $\lambda < 1$, the data clearly speak in favour of the Keane and Rogerson (2012) argument, according to which labor supply elasticity in the aggregate should fall within the interval of values ranging between 1 and 2. That is, our posterior density exhibits a median around 1.5, with the probability mass mostly concentrated within the [1,2] interval. Moving to γ, the prior density is rather loose while the posterior density narrows around values between 2 and 3, with a median estimate (2.4) essentially identical to the one estimated by Weinberg (2000), but also not significantly different from Acemoglu et al. (2004), who estimated γ to be close to 3.

When it comes to the remaining coefficients (see Table 1 and Figure C.2 in the Appendix), the data speak in favour of a *non-negligible* long-run feedback of gender-specific shocks to aggregate macro variables. On the one hand, the posterior densities of the labor supply feedbacks tend to concentrate around values close to 0 for both GDP, wages and - to a lesser extent - employment. On the other hand, the posterior estimates of the labor demand feedbacks shift away from values close to the 0-neighborhood, in particular for the case of GDP. Accordingly, we expect gender-specific labor demand to play a relatively more important role than labor supply. In addition, the coefficient measuring the effect of automation on employment (i.e., ν_{32}) is rather large and quite precisely estimated.[7]

Finally, one fine outcome of our identification procedure lies on the fact that the broad uniform sets enable us to remain fully agnostic on the *strength* of the long-run effects and, therefore, let the data speak. The rather narrow 90% credibility bands in Table 1 suggest, in fact, that the data embed valuable information on the parameters, thereby excluding the possibility that our posterior estimates are driven by the choice of the priors. The

[7] Recall that our identification procedure implies $\nu_{32} < 0$, by applying a minus in front of each draw. However, for consistency with the definition of the Gamma distribution, Table 1 reports the posterior estimates along a positive support.

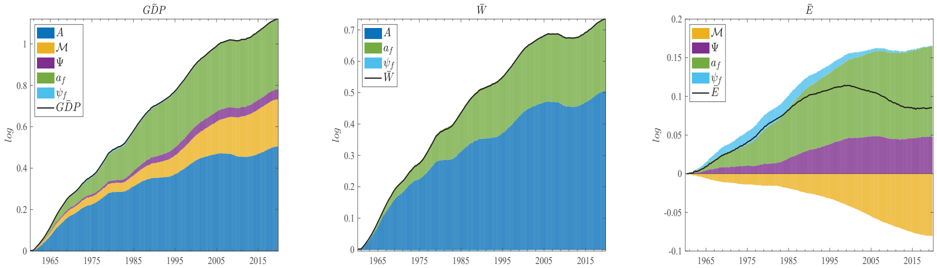

Figure 7: Contributions of the macro and female-specific *structural* trends to the trends in aggregate macro variables.

interpretation of the coefficients will be easier once we turn to the visual decomposition of the permanent components in the data into the five structural trends which is our next step.

Trends Decomposition. Let us now discuss the contribution of the different structural trends to each empirical trend in the model. Figure 6 decomposes the empirical trends of female employment gap and wage gap into gender-specific labor demand (green) and supply (light blue). These are the only shocks affecting gender differentials in the long-run. Three facts are worth noting. First, in the period 1960-1980, both positive labor demand and supply forces are in place. This explains the steep convergence of the employment gap, while the wage gap tends to stagnate over this period. Second, starting from roughly 1980 until the mid '90s, the dominant structural force driving the gender convergence in both employment and wage gaps is labor demand. Third, since the mid '90s, the gender convergence essentially stops and both employment and wage gaps exhibit a clear plateau lasting until today, implying a slow-down in the demand and supply forces that favored the gender convergence in the labor market.

The aim of this paper is to quantify the feedback effect of these gender-specific secular trends to the macroeconomy. Figure 7 presents an anatomy of the structural secular drivers for the US main macro aggregates.

The main result of this paper is that, consistently with the posterior estimates of the long-run elasticities in $\mathcal{V}$, gender-specific trends are important for the macroeconomy. This can be clearly seen by the large green area driving GDP, wage and employment trends upward. The gender-specific trends (green and light blue in right-hand side panel) are the sole driver of the positive trend in the US aggregate employment in the post-war period. Had it not been for the gender-specific trends, aggregate employment in the US would have stagnated.

In addition, the neutral technology trend A (blue area) is the main source of long-run co-movement between GDP and wages and is the main driver of trend US post-war economic growth. Perhaps surprisingly, aggregate labor supply Ψ (purple) does never play a significant role neither for GDP nor for employment: our model reads long-run changes in labor supply as mainly gender-specific. Finally, the trend in labor-displacing automation $\mathcal{M}$ (yellow) contributes significantly to GDP and employment, especially starting from approximately 1985. In particular, together with the slow-down of gender-specific trends, it is the main responsible of the negative trend experienced by aggregate employment from late '90s.

	$\Delta \bar{GDP}_t$				$\Delta \bar{E}_t$			
	Macro shocks	a_f	ψ_f	Total	Macro shocks	a_f	ψ_f	Total
1960-1969	2.00	0.54	0.05	2.6	-0.20	0.19	0.10	0.09
1970-1979	1.20	0.84	0.04	2.08	-0.09	0.29	0.09	0.29
1980-1989	0.88	1.01	-0.02	1.87	-0.28	0.35	-0.05	0.02
1990-1999	1.12	0.58	-0.03	1.67	-0.32	0.20	-0.06	-0.18
2000-2009	0.94	0.29	-0.02	1.21	-0.23	0.10	-0.05	-0.18
2010-2019	0.86	0.16	-0.01	1.01	-0.16	0.06	-0.02	-0.12

Table 2: Contribution of the macro and female-specific structural trends to the trend growth rates of real GDP and employment over time. Average of each corresponding decade.

One appealing feature of our empirical framework is that it accommodates for the possibility of performing a growth accounting exercise. Table 2 indeed reports the key result from the baseline specification. In the period 1970-1990, gender-specific shocks are responsible on average of approximately 1pp of the annualised GDP trend growth rate, which is about 50% of total growth. Furthermore, the significant contribution of gender-specific shocks is entirely attributable to labor demand factors, with labor supply playing no role. Yet, the 1pp loss in trend growth in the last twenty years is mainly due to the break in the gender convergence. As a matter of fact, the contribution of aggregate trends remained stable around 1pp over the last 50 years, while gender-specific trends plateaued precisely when GDP trend growth started its decline.

The second key results from table 2 refer to the drivers of the employment trend growth rate. In this case, at least for the first 20 years of the sample - i.e.: 1960-1979 -, both gender-specific labor demand and supply positively contribute to employment. Starting from the ’80s, however, labor supply fades away and the main contribution comes from labor demand. Aggregate shocks play almost no role until early ’90s, when the automation trend picks-up and together with the break in the gender convergence drag employment trend growth rate downward.

Overall, there are three main takeaways from the baseline specification. First, gender-specific trends are important for the macroeconomy. Second, according to our model, the slow-down in US economic growth in the last 25 years can be mainly attributed to the stagnating gender-convergence in the labor market. Last but not least, the Gender Revolution is mainly - if not entirely - a matter of labor demand factors.

However, in the light of the massive increase in female labor force participation over the last 50 years, this last statement is a hard bite to digest and arguing that gender-specific labor supply shocks do not play any significant role would probably be a hasty conclusion. In the interest of parsimony our baseline model might have potentially omitted some relevant features in the data, that once taken into account may change the overall picture. Data over the second half of the last century, for instance, not only document a massive entrance of women in the labor force, but especially an ever-increasingly number of highly educated (skilled) women entering the labor market. Now, suppose that the increase in women labor force participation is asymmetric, in the sense, that it is mainly

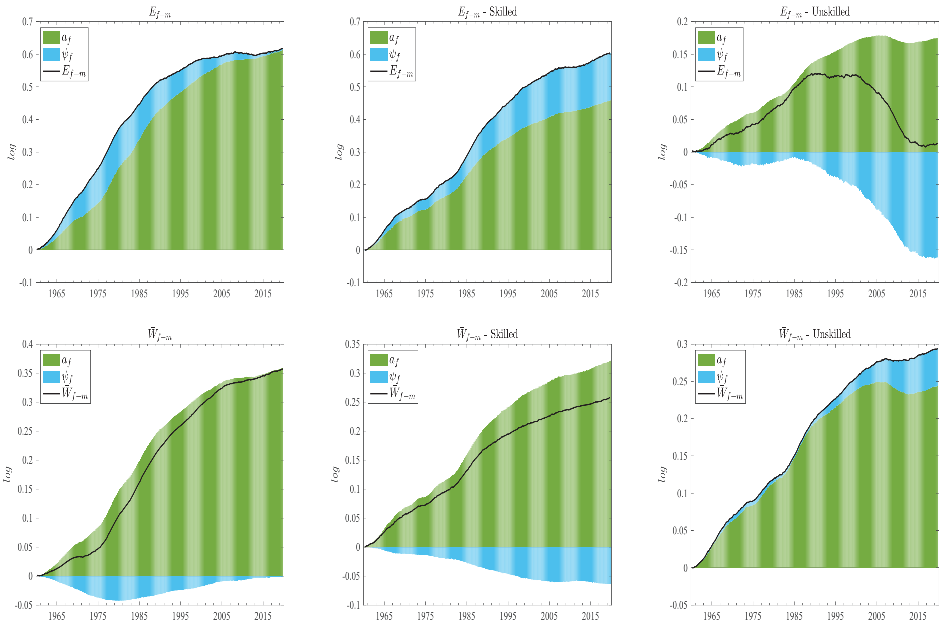

Figure 8: Contributions of the female-specific *structural* trends to the trends in female employment and wage gaps. Left panels: all workers (baseline). Middle panels: skilled workers. Right panels: unskilled workers.

induced by women who accessed higher levels of education. Then, this "asymmetric" shock may trigger a composition effect eventually leading to an increase in the average wage of women, as the share of highly educated workers in the women labor force has increased, ceteris paribus. Since the baseline specification does not control for the *skill dimension*, it would commingle this gender-specific *skilled* labor supply shock with labor demand factors, leading us to the wrong conclusion that labor demand factors are the *only* drivers of the Gender Revolution. In the following subsection, we investigate the implications of our empirical model when controlling for the *skill dimension* of the labor force. Concretely, based on the CPS dataset used by Dolado et al. (2021), we estimate two alternative specifications to the baseline, namely (i) a model with gender differential for *skilled* workers and (ii) a model with gender differentials for *unskilled* workers. Then, we set them side by side with our baseline specification and draw conclusions.

6 Accounting for the Skill and the Sectoral Dimension

In this section we investigate whether the important role of the gender convergence for trend US growth is confirmed once we take into account that the skill-mix and the size of the service sector in the US economy have both changed substantially.

6.1 The Role of Skills in the Gender Revolution

With respect to the baseline model, these alternative specifications differ only in the two gender variables embedded in the vector of endogenous variables Y_t in the SVAR. More precisely, we estimate two versions of our model in which the *aggregate* female employment and wage gaps are now replaced by the female employment and wage gaps of (i) *skilled* and (ii) *unskilled* workers respectively.

All the prior assumptions, the sample period and the estimation steps remain the same as in the baseline. In particular, we find that the uniform priors for the feedbacks of the gender-specific shocks to macro variables are reasonable also in the context of these alternative specifications. In addition, regarding the prior densities for λ and γ, there exists very scarce evidence documenting how these elasticities change when focusing on skilled/unskilled workers only. To the best of our knowledge, Acemoglu et al. (2004) is the only study documenting some heterogeneity between skilled and unskilled workers about the labor demand elasticity γ, though not about labor supply elasticity. According to the authors, the elasticity of substitution between skilled men and women should be higher - ranging between 4 and 10 - than in the case of unskilled men and women - spanning between 2.5 and 4 -. It is important to stress, nonetheless, that this study is based on data referring to the period 1940-60, which makes it unsuitable for the formulation of our priors, given that the degree of substitutability back then was likely to be very different from subsequent decades. In their multi-sector equilibrium model, Olivetti and Petrongolo (2014), instead, propose different calibrations for the labor demand elasticity within skilled workers - i.e.: $\gamma_S = 3.5$; $\gamma_S = 5$ - and unskilled workers - i.e.: $\gamma_U = 1.5$; $\gamma_U = 2.5$ -, while assuming no heterogeneity in the labor supply elasticity λ. Provided that the calibrations proposed by Olivetti and Petrongolo (2014) fall under the bell of the prior densities for both parameters, we finally decide to not change them.

First, compared to the aggregate, focusing only on skilled workers allows us to detect a *positive* gender-specific labor supply shock, which significantly contributes to the steep convergence in employment levels (upper-central panel in Figure 8) since the early '80s and at the same time responsible for the slow convergence in wages (lower-central panel).

Second, figure 9 does not support the dominant role of gender-specific labor demand over supply trend shock that we found in the baseline model. Here both gender-specific shocks are relevant for the aggregate macro variables. In particular, the gender-specific labor supply shock contribution to GDP (light blue) is no longer negligible. The contribution of gender-specific labor supply shocks is even relatively more important than labor demand when looking at the contribution to aggregate employment (bottom-central panel). Notice, in addition, that the light blue area has absorbed part of the contribution explained by the gender-specific labor demand (green), while the aggregate shocks - i.e.: labor supply (purple) and automation (yellow) - explains essentially the same as in the baseline. This confirms our concern of missing an important shock to the supply of skill that is undetectable when using aggregate data.

Let us now discuss the results from the specification controlling for *unskilled* workforce. the upper-right panel of figure 8 shows that the gender convergence in unskilled employment has been much more weaker than for skilled labor. By comparing the magnitudes on the y-axis, the convergence peak of unskilled employment - that stopped in 1990 - is comparable to the employment convergence reached by female skilled workers already in the '70s. In addition, starting from 1990 a strong negative gender-specific labor

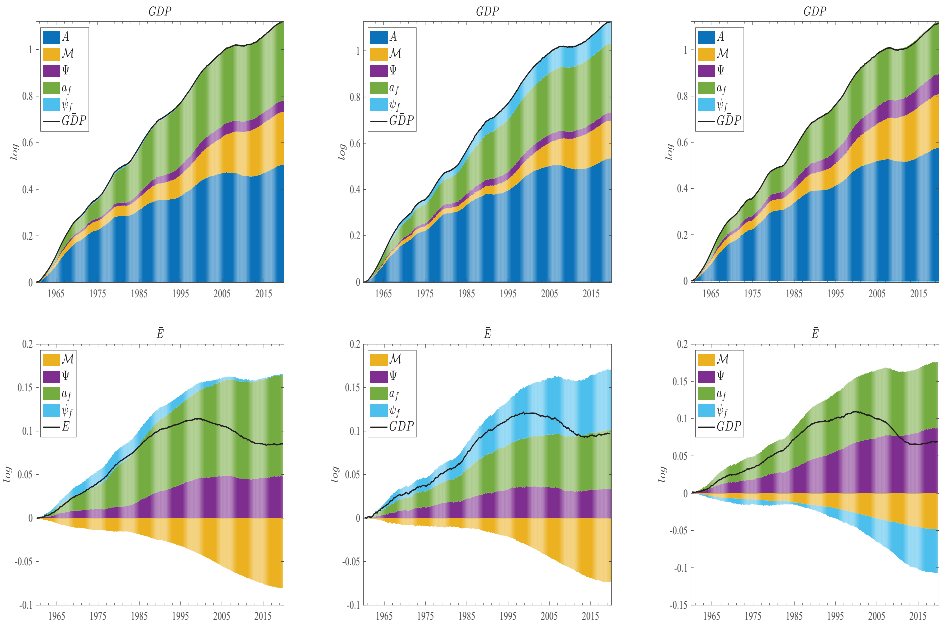

Figure 9: Contributions of the macro and female-specific *structural* trends to the trends in aggregate macro variables. Left panels: all workers (baseline). Middle panels: skilled workers. Right panels: unskilled workers.

supply shock dragged *unskilled* employment gap down until the end of the sample to levels of gender inequality similar to those observable in the 60s. Yet, this forceful negative labor supply shocks produced a steeper convergence in the wage rates in the final part of the sample, when the gender-specific labor demand substantially slowed down.

Figure 9 shows the trend decomposition of aggregate variables in its right column. Similarly to the results in the baseline model, the (*unskilled*) gender-specific labor demand trend dominates the contribution of gender specific labor supply. Only labor demand factors matter for the secular dynamics of GDP and no effects whatsoever arising from labor supply factors. In contrast with the baseline model, in which labor supply factors played some role in driving the upward trend of aggregate employment, the *unskilled* gender-specific labor supply contributes to the slowdown of the employment trend (right panel) over the entire sample and plays a non-negligible role especially from the '90s.

The key result from this section is that the Gender Revolution is not driven only by gender-specific demand factors. As a matter of fact, the results from the alternative specifications show that, once controlling for the skill dimension of the workforce, gender-specific labor supply factors do play a non-negligible role, in driving both gender differentials and macroeconomic outcomes. The reason why labor supply factors disappear in the aggregate, paving the way to the hegemony of labor demand factors, is twofold: (i) there exists a *positive* labor supply shock to *skilled* workers that turns out to be commingled with the labor demand shock in the baseline model, leading us to hold only labor demand accountable for the macroeconomic effects of the Gender Revolution; (ii) a forceful *negative* labor supply shock to *unskilled* workers partially off-sets the effects of the labor supply shock to *skilled* workers, partially canceling out the contribution of labor supply factors to the gender differentials in the aggregate.

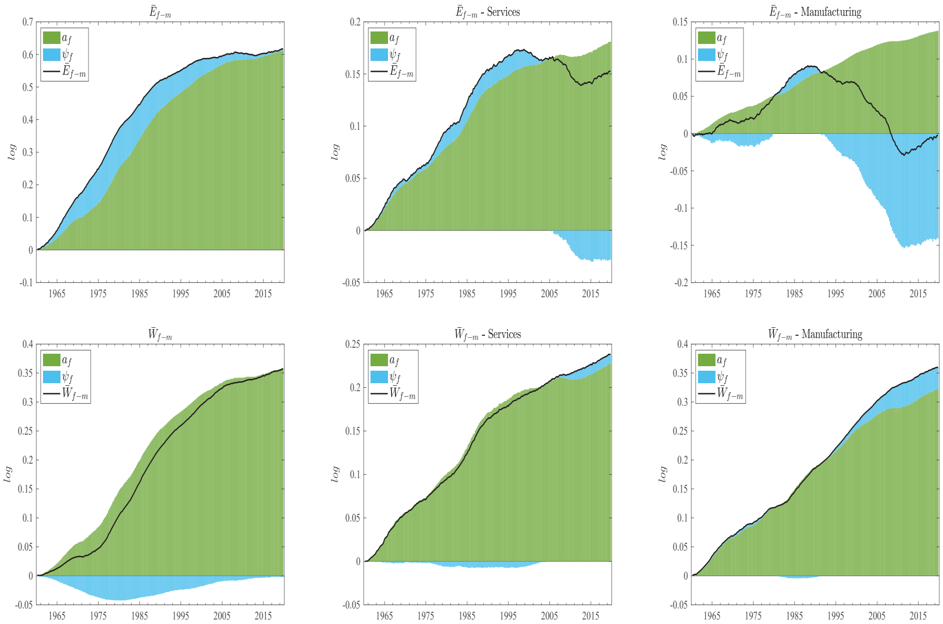

Figure 10: Contributions of the female-specific *structural* trends to the trends in female employment and wage gaps. Left panels: all workers (baseline). Middle panels: workers employed in services. Right panels: workers employed in manufacturing.

6.2 The Role of Sectors in the Gender Revolution

In figure 10 we compare the gender convergence in the aggregate (left column) against the gender convergence in the service sector (central column) and the gender convergence in the manufacturing sector (right column). We remark that there gender convergence also within both sectors. Therefore, the aggregate gender convergence is not an artifact of the service sector (more female intensive) becoming bigger over time. In fact, the employment convergence is steeper in the service sector than in the manufacturing sector while the wage convergence is slightly steeper in the manufacturing sector. To fit these dynamics, the model needs a negative female-specific labor supply shock in the manufacturing labor supply to a declining convergence in employment with a rather steep convergence in wages. in that sense, the dynamics observed in the manufacturing sector are reminiscent of our results for the convergence among low-skilled workers. The aggregate dynamics are somewhat similar to the dynamics in the service sector which is largely bigger than the manufacturing sector, especially toward the end of the sample. It is also important to keep in mind that the sum of employment in services and manufacturing is far from being equivalent to total employment.

When it comes to aggregate dynamics (plotted in figure 11, we remark that gender-specific factors are important for aggregate output and employment also when we estimate our sectoral models.

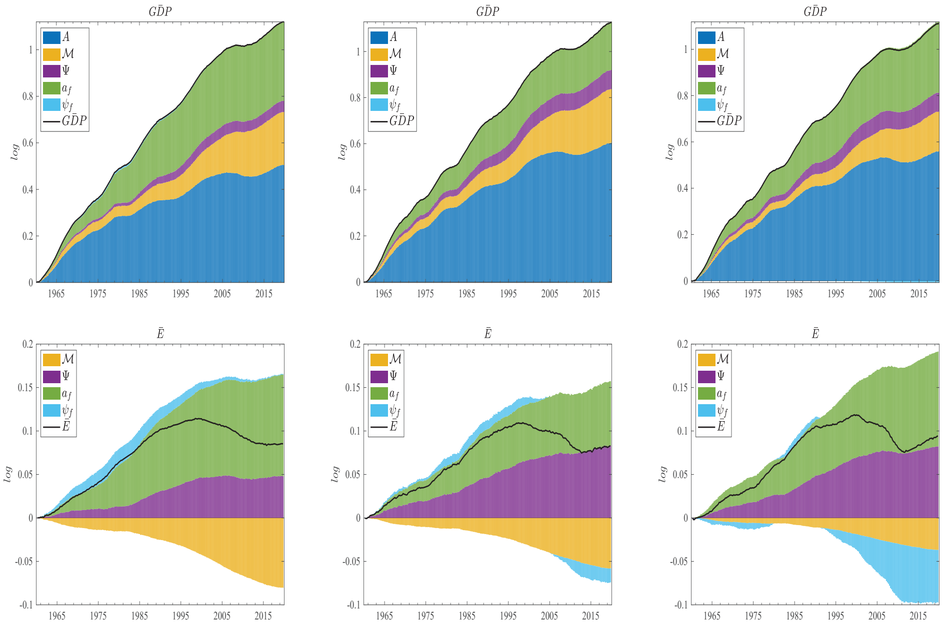

Figure 11: Contributions of the macro and female-specific *structural* trends to the trends in aggregate macro variables. Left panels: all workers (baseline). Middle panels: workers employed in services. Right panels: workers employed in manufacturing.

7 Conclusions

In this paper, we investigate the macroeconomic implications of the Gender Revolution. We quantify its impact on post-war economic growth in the US in terms of GDP, employment and wages and we shed light on the factors behind the secular convergence of employment and wages between females and males. To address our research question, we estimate a SVAR with common trends à la Del Negro et al. (2017) and Crump et al. (2019) and propose a decomposition of the empirical (reduced-form) trends into selected unobserved structural trends, that is motivated by economic theory.

Our first contribution is methodological. We propose explicitly mapping the estimated permanent component and five structural trends using economic theory as prior information to identify the model. This represents a key distinction from previous studies estimating VARs with stochastic trends. Note that our methodology can be applied to countless questions related to secular trends in the data and to investigate the link between growth and inequality.

Second, our empirical model documents the importance of gender-specific structural forces not only for the reduction of gender inequality (gender convergence) in the labor market, but also for economic growth. In particular, we show that gender-specific slow-moving trends account for up to 50% of the GDP trend growth rate over the period 1960-1990. Furthermore, according to our model, the flattening of the gender convergence started in the '90s is accountable for the marked slow-down observed in trend growth over the last 25 years. This result implies that policies aimed at completing the gender convergence process may deliver high payoffs in terms of economic growth.

The third contribution of the paper is to quantify the individual role of the gender-specific labor demand and supply in driving the secular convergence of the gender differentials. From the estimation of our baseline model, labor demand shocks completely overshadow labor supply. Hence, according to the baseline results, one would conclude that labor supply plays no role whatsoever in explaining the gender convergence. However, this conclusion turns out to be premature once controlling for the skill dimension of the labor market. Exploiting more disaggregated data we are able to rationalise the main results of the baseline model. More specifically, the macroeconomic implications of the Gender Revolution are not solely due to labor demand, but also to (*skilled*) labor supply. The baseline model fails to identify an important shock to the supply of skills, because the *skilled* labor supply trend is commingled with the labor demand. Yet, the presence of two opposite labor supply shocks to skilled and unskilled workers cancels out in the aggregate eventually downplaying the role of labor supply in driving the secular convergence of gender differentials in the US labor market.

Finally, we are currently working on an important extension of our model. In the current version of our analysis, the gender convergence is driven by female-biased shocks. However, male-biased shocks are also potentially important. In particular, one can conjecture that the decline in male labor force participation is driven in part also by male-specific factors. We plan to shed light on this question over the coming weeks. In addition, our framework can be used also to study gender differences at business cycle frequencies (cf. Albanesi and Şahin (2018)). We plan to investigate this question in future research.

References

Acemoglu, D., D. H. Autor, and D. Lyle (2004, June). Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury. Journal of Political Economy 112(3), 497–551.

Acemoglu, D. and P. Restrepo (2020). Robots and Jobs: Evidence from US Labor Markets. Journal of Political Economy 128(6), 2188–2244.

Albanesi, S. (2019). Changing business cycles: The role of women's employment. Technical report, National Bureau of Economic Research.

Albanesi, S. and C. Olivetti (2016). Gender roles and medical progress. Journal of Political Economy 124(3), 650–695.

Albanesi, S. and A. Şahin (2018). The gender unemployment gap. Review of Economic Dynamics 30, 47–67.

Ascari, G. and L. Fosso (2021, December). The inflation rate disconnect puzzle: On the international component of trend inflation and the flattening of the phillips curve. Working paper no. 733, De Nederlandsche Bank.

Attanasio, O., H. Low, and V. Sánchez-Marcos (2008). Explaining changes in female labor supply in a life-cycle model. American Economic Review 98(4), 1517–52.

Bergholt, D., F. Furlanetto, and N. Maffei-Faccioli (2022). The decline of the labor share: new empirical evidence. American Economic Journal: Macroeconomics 14(3), 163–198.

Bianchi, F., G. Nicolò, and D. Song (2023). Inflation and real activity over the business cycle. Technical report, National Bureau of Economic Research.

Blundell, R. and T. Macurdy (1999). Labor supply: A review of alternative approaches. In O. Ashenfelter and D. Card (Eds.), Handbook of Labor Economics, Volume 3 of Handbook of Labor Economics, Chapter 27, pp. 1559–1695. Elsevier.

Buera, F. J. and J. P. Kaboski (2012). The rise of the service economy. American Economic Review 102(6), 2540–69.

Coibion, O., Y. Gorodnichenko, and M. Ulate (2018). The cyclical sensitivity in estimates of potential output. Brookings Papers on Economic Activity.

Crump, R. K., S. Eusepi, M. Giannoni, and A. Sahin (2019). A unified approach to measuring u. Brookings Papers on Economic Activity, 143–239.

Del Negro, M., D. Giannone, M. Giannoni, and A. Tambalotti (2017, May). Safety, liquidity, and the natural rate of interest. Staff Reports 812, Federal Reserve Bank of New York.

Dolado, J. J., G. Motyovszki, and E. Pappa (2021). Monetary policy and inequality under labor market frictions and capital-skill complementarity. American economic journal: macroeconomics 13(2), 292–332.

Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. Biometrika 89(3), 603–615.

Fernández, R., A. Fogli, and C. Olivetti (2004). Mothers and sons: Preference formation and female labor force dynamics. The Quarterly Journal of Economics 119(4), 1249–1299.

Fukui, M., E. Nakamura, and J. Steinsson (2023). Women, wealth effects, and slow recoveries. American Economic Journal: Macroeconomics.

Furlanetto, F., A. Lepetit, Ørjan Robstad, J. Rubio-Ramírez, and P. Ulvedal (2023). Estimating Hysteresis Effects. forthcoming, American Economic Journal: Macroeconomics.

Galor, O. and D. N. Weil (1996). The gender gap, fertility, and growth. The American Economic Review, 374–387.

Goldin, C. (2006). The quiet revolution that transformed women's employment, education, and family. American economic review 96(2), 1–21.

Goldin, C. and L. F. Katz (2002). The power of the pill: Oral contraceptives and women's career and marriage decisions. Journal of political Economy 110(4), 730–770.

Greenwood, J., A. Seshadri, and M. Yorukoglu (2005). Engines of liberation. The Review of Economic Studies 72(1), 109–133.

Hamermesh, D. S. (1993). Labor demand. Princeton University Press, Princeton, New Jersey.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). The macroeconomics of the quiet revolution: Understanding the implications of the rise in women's participation for economic growth and inequality. Research in Economics 71(3), 521–539.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The allocation of talent and us economic growth. Econometrica 87(5), 1439–1474.

Johnson, M. and M. Keane (2013). A dynamic equilibrium model of the us wage structure, 1968-1996. Journal of Labor Economics 31(1), 1 – 49.

Keane, M. and R. Rogerson (2012). Micro and macro labor supply elasticities: A reassessment of conventional wisdom. Journal of Economic Literature 50(2), 464–476.

Maffei-Faccioli, N. (2021). Identifying the sources of the slowdown in growth: Demand vs. supply. Technical report, Norges Bank Working Paper 9/2021.

Olivetti, C. and B. Petrongolo (2014, October). Gender gaps across countries and skills: Demand, supply and the industry structure. Review of Economic Dynamics 17(4), 842–859.

Taniguchi, H. and K. Yamada (2020). The race between technological progress and female advancement: Changes in gender and skill premia in oecd countries. arXiv preprint arXiv:2005.12600.

Weinberg, B. A. (2000). Computer use and the demand for female workers. Industrial and Labor Relations Review 53(2), 290–308.

Appendix

A A BVAR with Common Trends

The model presented in eq. (1)–(4) of Section 2 aims at decomposing the empirical trends $\bar{Y}_t$ into three aggregate macro drivers: (i) technology denoted by A_t; (ii) automation $\mathcal{M}_t$; (iii) labor supply Ψ_t. Technology and labor supply trends are the most natural drivers of growth in the economy. The automation trend is meant to capture those slow-moving factors contributing to the secular decline of labor share - see e.g.: Acemoglu and Restrepo (2020), Bergholt et al. (2022). In addition, the long-run behavior of employment and wage gender ratios is driven by two forces: (iv) gender-specific labor demand $a_{f,t}$; (v) gender-specific labor supply $\psi_{f,t}$. These are meant to capture any *positive* and *negative* co-movement between employment and wage gap, respectively. The distinction between aggregate and gender-specific drivers reminds respectively the labels "background forces" and "primary forces" used by Heathcote et al. (2017) in the context of a calibrated structural model used to study the drivers of economic growth in the US. Therefore, X_t is a 5×1 vector, so the number of observables n is equal to the number of structural trends q. Given the system of observables Y_t, the main challenge is to find a unique solution that allows us to reconcile the reduced-form objects in $\bar{Y}_t$ with the structural ones stacked in X_t.

A.1 Prior Assumptions

The initial conditions of the structural trends are distributed according to $X_0 \sim \mathcal{N}(\underline{X_0}, I_q)$. In principle, we do not have information about $\underline{X_0}$. However, one can use the information on $\underline{\bar{Y}_0}$ - the initial conditions of the empirical trends[8] - as well as on the prior coefficients in $\mathcal{V}$. Then, one can retrieve $\underline{X_0}$ by solving the system in eq. (2) - provided that the number of structural trends $q = n$, as it is the case in our model.[9] The initial conditions of the cycles are distributed according to $\hat{Y}_0 \sim \mathcal{N}(0_n, I_n)$. This assumption implies that cycles fluctuate symmetrically around a zero mean. Finally, the priors for the remainder model's coefficients are distributed according to:

$$\Sigma_u \sim \mathcal{IW}(\kappa_u, (\kappa_u + n + 1)\underline{\Sigma_u}) \tag{A.1}$$

$$\Sigma_e \sim \mathcal{IW}(\kappa_e, (\kappa_u + n + 1)\underline{\Sigma_e}) \tag{A.2}$$

$$\tilde{\Phi}|\Sigma_e \sim \mathcal{N}(\underline{\tilde{\Phi}}, \Sigma_e \otimes \underline{\Omega})\mathcal{I}(\tilde{\Phi}), \tag{A.3}$$

where $\tilde{\Phi} = vec(\Phi)$ and $\mathcal{I}(\tilde{\Phi})$ is an indicator function that is equal to one, when the VAR of the cycle block is stationary, zero otherwise. $\mathcal{IW}$ is the Inverse-Wishart distribution with κ degrees of freedom and mode $\underline{\Sigma}$. We assume the prior mode of trend shocks $\underline{\Sigma_u}$ to be diagonal and its non-zero elements are retrieved in the same fashion as we discussed for the initial conditions of the structural trends $\underline{X_0}$. We impose a rather tight prior around the prior mode, by setting the degrees of freedom $\kappa_u = 100$. Moving to the cycle block,

[8] Specifically, we set $\underline{\bar{Y}_0}$ equal to the average of the HP-filter trend growth rate from the pre-sample data.

[9] In section 7 of the appendix, we show how to derive priors for the structural trends by solving the implied system in eq. (2).

the priors for the lag coefficients are standard Minnesota with an overall tightness hyperparameter equal to 0.2 and the own-lag hyperparameters centered around zero, instead of one, since we are dealing with the stationary block. The prior mode of the transitory innovations $\underline{\Sigma_e}$ is assumed to be an identity matrix and rather uninformative, as the degrees of freedom are $\kappa_e = n + 2$.

One non-trivial task is to come up with reasonable priors for the elements of $\sigma^2_{\mathbf{u}} = [\sigma^2_A \;\; \sigma^2_{\mathcal{M}} \;\; \sigma^2_{\Psi} \;\; \sigma^2_{a_f} \;\; \sigma^2_{\psi_f}]'$, the vector stacking the shocks' volatilities of the *structural* trends in X_t. The reason is because the structural trends are *unobservable* in the first place. Nonetheless, it is still possible to form fairly non-judgmental priors on these structural volatilities by combining two pieces of information we already possess, namely: (i) the data and (ii) the theory-based prior beliefs on the free parameters in $\mathcal{V}(\nu)$. To see this, recall that empirical and structural trends are linked by the linear relationship $\bar{Y}_t = \mathcal{V}X_t$ and that $X_t = c + X_{t-1} + u_t$. Without loss of generality, one can express the empirical trends in their growth rates, as follows:

$$\Delta\bar{Y}_t = \mathcal{V}(c + u_t)$$

This equation implies that the covariance matrix of the empirical trends in growth rates is denoted by $\Sigma_{\Delta\bar{Y}} = \mathcal{V}'\Sigma_u\mathcal{V}$. Then, provided that the covariance matrix Σ_u is diagonal, the following linear relations apply:

$$\sigma^2_{G\bar{D}P} = \sigma^2_A + \sigma^2_{\mathcal{M}} + \sigma^2_{\Psi} + \nu^2_{14}\sigma^2_{a_f} + \nu^2_{15}\sigma^2_{\psi_f} \qquad \sigma^2_{\bar{W}} = \sigma^2_A + \nu^2_{24}\sigma^2_{a_f} + \nu^2_{25}\sigma^2_{\psi_f}$$

$$\sigma^2_{\bar{E}} = \nu^2_{32}\sigma^2_{\mathcal{M}} + \sigma^2_{\Psi} + \nu^2_{34}\sigma^2_{a_f} + \nu^2_{35}\sigma^2_{\psi_f} \qquad \sigma^2_{\bar{E}_{f-m}} = \nu^2_{44}\sigma^2_{a_f} + \nu^2_{45}\sigma^2_{\psi_f}$$

$$\sigma^2_{\bar{W}_{f-m}} = \sigma^2_{a_f} + \sigma^2_{\psi_f}$$

On the left-hand side of each equation, there are the volatilities of the empirical trends in growth rates, while on the right-hand side, there are the coefficients of $\mathcal{V}$ and volatilities of the structural shocks. The empirical volatilities are available in the data and the parameters ν_{ii} are simply the values around which the prior density of the long-run elasticities is centered. The only *unknowns* are the structural volatilities. It turns out that is straightforward to retrieve the structural volatilities in $\sigma^2_{\mathbf{u}}$, as they are the unknowns of a linear system of 5 equations in 5 unknowns and, therefore, there always exists a unique solution to the system. Consistently, this is how we proceed in practice. First, back out the empirical volatilities from the HP-filter trend growth rates of the endogenous variables using pre-sample training. Second, plug the empirical volatilities and the prior means of the parameters in $\mathcal{V}$. Solve the system for the unknown volatilities and use them to center the prior density of the structural volatilities.

Finally, notice that the very same reasoning applies when forming priors for the initial conditions and the drifts of the structural trends. Accordingly, the initial conditions X_0 should be centered around $X_0 = \mathcal{V}\bar{Y}_0$, with $\bar{Y}_0$ being the *last period*'s empirical trend *in levels* (last period in the training sample). As for the drifts, the constants c should be centered around $c = \mathcal{V}\mathbb{E}(\Delta\bar{Y}_t)$, with $\mathbb{E}(\Delta\bar{Y}_t)$ being the average of the empirical trends *in growth rates* (in the training sample).

A.2 Estimation of the State Space with Gibbs Sampling

Consider the unobserved states of the model in section 2 in the following stacked formulation:

$$\begin{bmatrix} \mathcal{V}X_t \\ \hat{Y}_t \end{bmatrix} = \begin{bmatrix} \mathcal{V}c \\ 0 \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & A \end{bmatrix} \begin{bmatrix} \mathcal{V}X_{t-1} \\ \hat{Y}_{t-1} \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} \mathcal{V}u_t \\ e_t \end{bmatrix} \tag{A.4}$$

and the Covariance matrix of the model is given by Σ:

$$\Sigma = \begin{bmatrix} \mathcal{V}'\Sigma_u\mathcal{V} & 0 \\ 0 & \Sigma_e \end{bmatrix} \tag{A.5}$$

Then, the model samples 50000 draws and retains the last 10000 draws from a Gibbs algorithm, according to the following steps:

1. Draw from the joint distribution $X_{0:T}, \hat{Y}_{-p+1:T}, \nu \mid c, A, \Sigma_u, \Sigma_e, Y_{1:T}$, which is given by the product of the marginal posterior of ν - vector of free parameters in $\mathcal{V}$ - conditional on the other parameters $\nu \mid c, A, \Sigma_u, \Sigma_e, Y_{1:T}$ and the distribution of the unobserved states conditional on ν and the other parameters $X_{0:T}, \hat{Y}_{-p+1:T} \mid \nu, c, A, \Sigma_u, \Sigma_e, Y_{1:T}$.

    (a) $p(\nu \mid c, A, \Sigma_u, \Sigma_e, Y_{1:T}) \propto \mathcal{L}(Y_{1:T} \mid \nu, c, A, \Sigma_u, \Sigma_e)p(\nu)$,
    where $\mathcal{L}(Y_{1:T} \mid \nu, c, A, \Sigma_u, \Sigma_e)$ is the likelihood of the data obtained from the Kalman filter applied to the state space of the model. The posterior of ν does not have a known solution, therefore we approximate it by introducing a Metropolis-Hastings step.

    (b) Draws from p($X_{0:T}, \hat{Y}_{-p+1:T} \mid \nu, c, A, \Sigma_u, \Sigma_e, Y_{1:T}$) are obtained implementing Durbin and Koopman (2002) simulation smoothing algorithm.

2. Draw from the joint distribution $A, c, \Sigma_u, \Sigma_e \mid X_{0:T}, \hat{Y}_{-p+1:T}, Y_{1:T}$. The estimation of the remaining parameters is relatively straightforward, provided that the unobserved states follow rather standard vector autoregressive laws of motion.

    (a) **Trend Block.** the posterior distribution of Σ_u is given by:

    $$p(\Sigma_u \mid X_{0:T}) = \mathcal{IW}(\underline{\Sigma_u} + \underbrace{\sum_{t=1}^{T}(X_t - X_{t-1})(X_t - X_{t-1})'}_{S_u}, \kappa_u + T)$$

    The posterior distribution of the vector of drifts c conditional on the Σ_u and $X_{0:T}$ is obtained from a standard Normal.

    (b) **Cycle Block.** The posterior distributions of the lag coefficients in A and the covariance matrix Σ_e of the stationary VAR are standard:

    $$p(\Sigma_e \mid \hat{Y}_{0:T}) = \mathcal{IW}(\underline{\Sigma_e} + S_e, \kappa_e + T)$$
    $$p(A \mid \Sigma_e, \hat{Y}_{0:T}) = N\left(vec(\mathcal{A}), \Sigma_e \otimes \left(\sum_{t=1}^{T} \hat{Z}_t\hat{Z}_t' + \underline{\Omega}^{-1}\right)^{-1}\right)$$

where $\hat{Z}_t = (\hat{Y}'_{t-1}, \ldots, \hat{Y}'_{t-p})$,
$\mathcal{A} = \left(\sum_{t=1}^{T} \hat{Z}_t \hat{Z}'_t + \underline{\Omega}^{-1}\right)^{-1} \left(\sum_{t=1}^{T} \hat{Z}_t \hat{Y}'_t + \underline{\Omega}^{-1} \underline{A}\right)$,
$S_e = \sum_{t=1}^{T} e_t e'_t + (\mathcal{A} - \underline{A})' \underline{\Omega}^{-1} (\mathcal{A} - \underline{A})$

B DATA

Data available on the FRED website is listed in the table below along with their identification code.

DATA	CODE
Real Gross Domestic Product per capita	A939RX0Q048SBEA
Non-farm business sector: real compensation per hour	COMPRNFB
Employment level, thousands of persons	CE16OV
Population-level, thousands of persons	CNP16OV
Employment-to-Population ratio	EMRATIO
Women Employment-to-Population ratio	LNS12300002
Men Employment-to-Population ratio	LNS12300001
Women nominal weekly earnings	LES1252882700Q
Men nominal weekly earnings	LES1252881800Q

Table 3: US data definitions and identification codes

Transformations. Data only available at monthly frequency (e.g.: employment, population, etc.) are transformed into quarterly by taking the three-month average of each corresponding quarter. Real aggregate wages per capita are retrieved from the following product $COMPRNFB \times \frac{CE16OV}{CNP16OV}$. Regarding the gender employment and wage gaps, these are computed as the ratios of female-to-male employment and wage rates, respectively. Before computing the ratios, the gender-specific wage rates are transformed into hourly wage rates, dividing them by 52 (number of working weeks in a year). Also, unfortunately, gender-specific wage rates are only available from 1979Q1. To missing observations, the period spanning 1960-1978 is filled by the earnings data available from the Annual Social and Economic Supplements, Current Population Survey, U.S. Census Bureau. These data are at annual frequency, thus we decide to get the intra-annual observations by using standard interpolation techniques.

Gender-specific Data by Sectors and Skills. Following Dolado et al. (2021), we download data on hourly wages and employment by skills and sector from the publicly available Current Population Survey (CPS) produced by the United Census Bureau and the Bureau of Labor Statistics downloadable at `http://data.nber.org/morg/annual/`. Then, we edit the STATA code used by Dolado et al. (2021) to merge the CPS survey into a unified dataset and back out the gender employment and wage gaps by skills and sectors. We refer to the online Appendix of Dolado et al. (2021) for a detailed discussion on how the CPS dataset is merged. For our analysis, we are interested in retrieving the employment level and hourly wage rate of women and men by skills and sectors. Regarding the skills dimension, we limit to split individuals into those with at least *some college* experience (skilled) and those who do not have any college education at all (unskilled). Hence, we can retrieve the employment level and hourly wages of women and men skilled and unskilled workers.

Manufacturing	*Services*
Nonmetallic mineral products,	Wholesale trade
Primary metals and fabricated metal products	Retail trade
Machinery manufacturing	Transportation and warehousing
Computer and electronic products	Publishing industries
Electrical equipment, appliance manufacturing	Motion picture and sound recording industries
Transportation equipment manufacturing	Broadcasting
Wood products	Internet publishing and broadcasting
Furniture and fixtures manufacturing	Telecommunications
Miscellaneous and not specified manufacturing	Internet providers and data processing services
Food manufacturing	Other information services
Beverage and tobacco products	Finance, Insurance, Real Estate
Textile, apparel, and leather manufacturing	Rental and leasing services
Paper and printing	Professional and Technical services
Petroleum and coal products	Management of companies and enterprises
Chemical manufacturing	Administrative and support services
Plastics and rubber products	Waste management and remediation services
	Educational services
	Hospitals, Health care services
	Social assistance
	Food services and drinking places
	Private households

Table 4: Industries included in the manufacturing and services sectors.

C Additional Tables & Figures

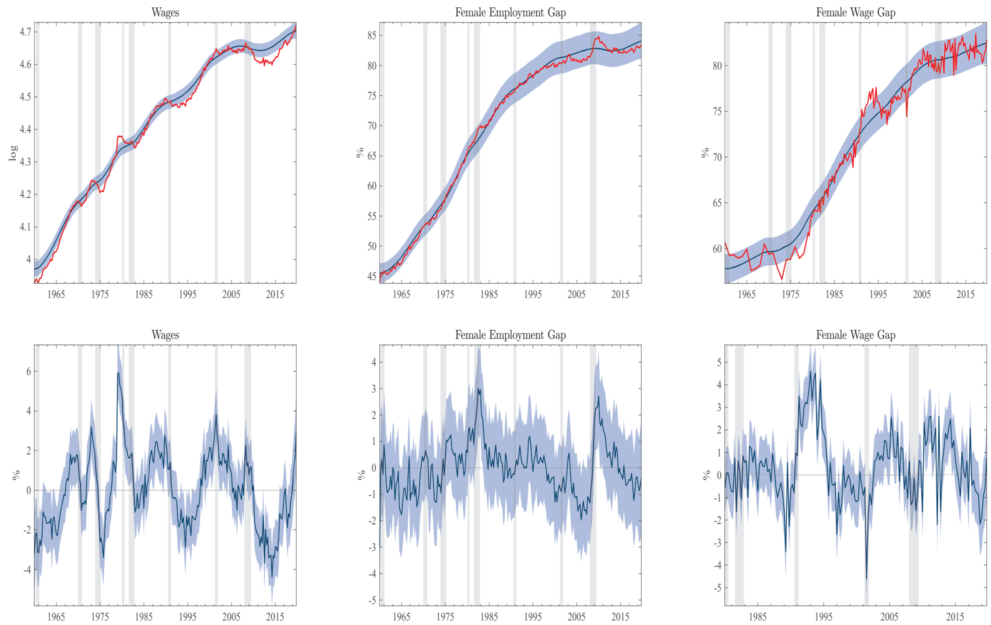

Figure C.1: Empirical trends and cycles, baseline model. Observed data (red, solid), median trend and cycle estimates (blue solid). 68% uncertainty band shaded area.

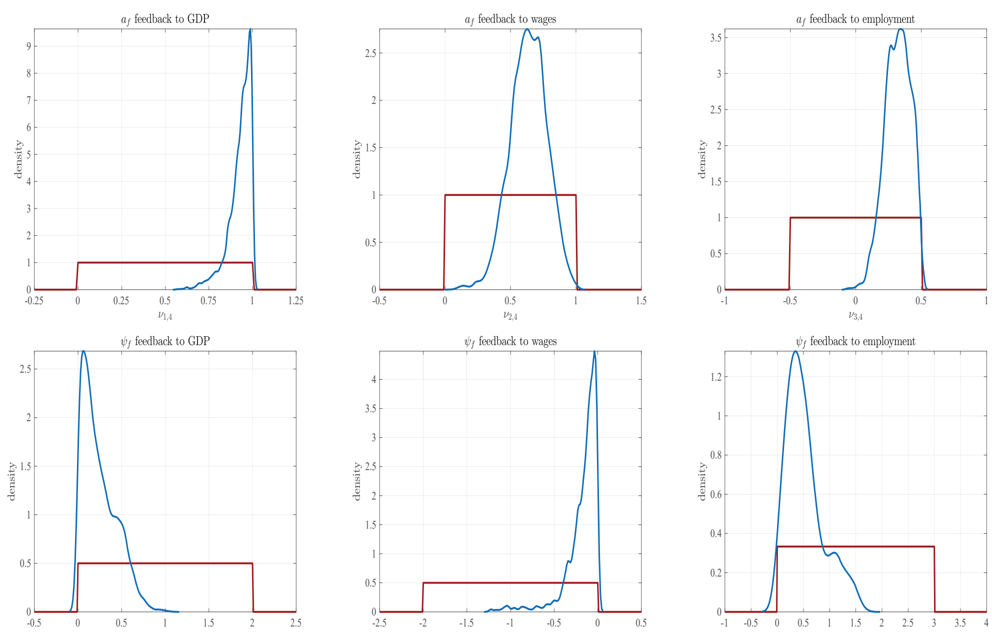

Figure C.2: Priors (red) and posteriors (blue) of the female-specific trends' feedbacks to the aggregate macro trends.

	$\Delta\bar{E}_{f,t}$			$\Delta\bar{W}_{f,t}$		
	a_f	ψ_f	Total	a_f	ψ_f	Total
1960-1969	0.97	0.67	1.64	0.57	-0.23	0.34
1970-1979	1.50	0.56	2.06	0.88	-0.19	0.69
1980-1989	1.83	-0.31	1.52	1.07	0.11	1.18
1990-1999	1.04	-0.38	0.66	0.61	0.13	0.74
2000-2009	0.52	-0.30	0.22	0.30	0.11	0.41
2010-2019	0.29	-0.16	0.13	0.17	0.06	0.23

Table 5: Contribution of the female-specific structural trends to the trend growth rates of female employment and wage gaps over time. Average of each corresponding decade.